Converged Networks and Fibre Channel over Ethernet

Table of contents

Abstract

In this paper we examine requirements for converged networking up to the server edge and beyond for end-to-end network environments. We want you to understand how the HP vision for converging technology, management tools, and partner product portfolios aligns with those requirements. We also compare HP and Cisco approaches to converged networking so you can see the likely outcomes of the two strategies. If you are a senior executive, CTO, CIO, IT manager, or business manager looking to deploy end-to-end converged networks in the next few years, this white paper can help you make more informed decisions about converged networking in your IT environment.

Introduction

Traditional data centers typically have underused capacity, inflexible purpose-built resources and high management costs. Infrastructure designs in such data centers include separate, heterogeneous network devices for different types of data. Each device adds complexity, cost, and management overhead. Many data centers support three or more types of networks and therefore require unique switches, network adapters, and network management systems to unify the networks.

Network convergence in the data center is a simpler and more economical solution than using separate purpose-built networks for data, management, and storage. Network convergence simplifies data center infrastructure by consolidating block-based storage and traditional IP-based data communications networks onto a single converged Ethernet network.

To achieve network-wide convergence, the IT industry has proposed and ratified the IEEE 802.1 Data Center Bridging (DCB) network protocol standards. The industry as a whole is deciding which implementation paths to pursue. In this paper we look at issues and choices IT organizations and vendors face when implementing these standards in networking products and data center infrastructures.

Converged networks

To better understand the current state of data center networking and the difficult tradeoffs facing IT organizations today, we should look at FCoE and DCB network standards. Protocol standards ratified by the IEEE 802.1 Data Center Bridging Task Group[1] affect network hardware, architecture, and behavior in converged end-to-end network infrastructure. Attempts to create end-to-end converged networks over the past decade have used the Fibre Channel Protocol, InfiniBand, and iSCSI. All have found their place in the market; but for both technological and business reasons, none have gained widespread acceptance as the Converged Network standard.

In compliance with ratified industry standards, HP has successfully implemented converged Fibre Channel over Ethernet (FCoE) networks at the server edge using top-of-rack (ToR) switches, Virtual Connect (VC) FlexFabric module, FlexFabric Adapters, and Converged Network Adapters (CNAs). The server edge is the first hop for networks in multi-hop or multi-switch network environments. The next major challenge is to extend that convergence throughout the data center network.

Figure 1 shows a single DCB-compliant network carrying IP, storage (FCoE), and cluster traffic using RDMA over Converged Enhanced Ethernet (RoCEE). You can read more about RoCEE and the expected trends in clustering in the "Cluster interconnects" section later in this paper.

Figure 1: Converged network, storage, and cluster traffic flows over 10GbE links.

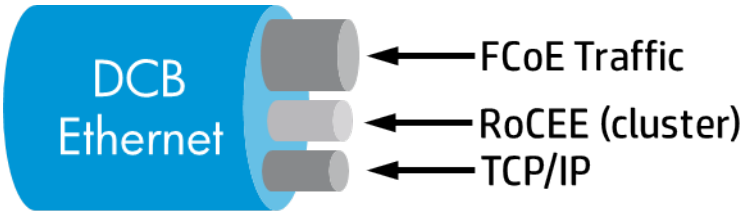

[1] You can get more information on the IEEE 802.1 DCB Task Group at http://www.ieee802.org/1/pages/dcbridges.html

The benefits of convergence are clear: Converged networking reduces the number of required I/O ports. This means that smaller servers (including blades) with fewer available option card slots are ideally suited for taking advantage of converged networks and the reduced requirement for I/O ports for full connectivity. Furthermore, the savings from not purchasing and operating several fabric types and their server connections is substantial. Access to any and all resources in the data center using simple wire-once servers reduces deployment, diagnosis, management, and operating costs.

Data Center Bridging

An informal consortium of network vendors originally defined a set of enhancements to Ethernet to provide enhanced traffic management and lossless operation. The consortium's proposals have become a standard from the Data Center Bridging (DCB) task group within the IEEE 802.1 Work Group.

The DCB standards define four new technologies:

- Priority-based Flow Control (PFC): 802.1Qbb allows the network to provide link level flow control (priority based pause) for different traffic classes.
- Enhanced Transmission Selection (ETS): 802.1Qaz defines the scheduling behavior of multiple traffic classes, including strict priority and minimum guaranteed bandwidth capabilities. This should enable fair sharing of the link, better performance, and metering.
- Data Center Bridging Exchange Protocol (DCBX): 802.1Qaz supports discovery and configuration of network devices that support PFC and ETS parameter negotiation between link partners. DCBX is specified as part of the same standard as ETS.
- Quantized Congestion Notification (QCN): 802.1Qau supports end-to-end flow control in a switched LAN infrastructure and helps eliminate sustained, heavy congestion in an Ethernet fabric. Before the network can use QCN, you must implement QCN in all components in the DCB data path (CNAs, switches, and so on). QCN enabled networks work with PFC to avoid dropping packets and ensure a lossless environment.

The DCB task group completed standards ratification in 2011.

Fibre Channel over Ethernet (FCoE)

FCoE is a protocol that wraps Fibre Channel frames in Ethernet frames for transporting them over Ethernet links. FCoE is widely accepted as an enabler of converged I/O in servers. FCoE involves minimal changes to the Fibre Channel protocol. FCoE packets encapsulate Fibre Channel frames inside Ethernet frames (Figure 2).

Figure 2: The FCoE protocol embeds FC frames within Ethernet frames.

FCoE on a DCB network mimics the lightweight nature of native Fibre Channel protocols and media. It does not incorporate TCP or even IP protocols. This means that FCoE is a layer 2 (non-routable over IP) protocol just like Fibre Channel. FCoE is primarily focused on local communication within a data center. The main advantage of FCoE is that switch vendors can easily implement logic for converting FCoE on a DCB network (FCoE/DCB) to native Fibre Channel in high-performance switch silicon.

In a single-hop FCoE/DCB network architecture, a function within a switch known as a Fibre Channel Forwarder (FCF) passes encapsulated Fibre Channel frames between a servers' CNA and the Fibre Channel storage area networks (SAN) where the Fibre Channel storage targets are connected. An FCF is typically an Ethernet switch with DCB, legacy Ethernet, and legacy Fibre Channel ports. Examples of FCFs include HP VC FlexFabric modules and HP Networking 5820X top-of-rack access switches with Fibre Channel option modules.

FCoE has several advantages:

- The CNA appears to the OS as a FC HBA and a NIC and uses existing OS device drivers. FCoE is implemented internal to the CNA and its device driver, but the fact that FCoE is in use is not presented to the server OS.
- Uses the existing Fibre Channel security and management model
- Makes storage targets that are provisioned and managed on a native Fibre Channel SAN transparently accessible through an FCoE FCF

However, there are also some challenges with FCoE:

- For anything other than a single-hop scenario (CNA to first switch where storage traffic is broken out to native Fibre Channel uplinks), it must be deployed on a DCB-enabled Ethernet network, as a lossless transport is needed.
- Requires CNAs and new DCB-enabled Ethernet switches between the servers and FCFs (to accommodate DCB)
- Is a non-routable over-IP protocol primarily used within the data center, as are native Fibre Channel protocols today
- Requires an FCF device to connect the DCB network to the legacy Fibre Channel SANs and storage
- Requires validating a new fabric infrastructure that converges LAN communications and Fibre Channel traffic over DCB-enabled Ethernet. Validating the network ensures that you have applied proper traffic class parameters to meet your IT organizations' business objectives and service level agreements.

The first step of server converged I/O at the server edge (network's first hop) is available now. Organizations achieve significant equipment cost savings by reducing the number of interfaces required for each server and the number of interconnects providing that first hop. Extending converged network as an end-to-end implementation in multi-hop networks holds a different set of challenges.

Challenges for end-to-end network convergence

This section describes challenges to network convergence that arise both from the varied demands placed on network infrastructure in enterprise environments, and from established network technologies that offer alternatives to FCoE

Many people within the IT community expect FCoE to allow end-to-end use of Ethernet for storage as well as traditional IP network traffic. The networking community expects FCoE to be an open, interoperable protocol where FCoE and TCP traffic co-exist on multivendor Layer 2 networks. The storage community expects everything currently running on Fibre Channel to operate over Ethernet. They also expect to include the installed base of Fibre Channel hardware in that network for its useful lifetime. CFOs expect to reduce maintenance costs by ultimately merging the physical Fibre Channel and Ethernet infrastructure.

Vendors are currently implementing DCB standards in network hardware and firmware. HP Virtual Connect FlexFabric technology has provided standards-based converged networking to the server edge since 2010. Realistically, however, FCoE can't be all things to all parties, at least not at the same time.

Cisco is one of the few vendors promoting end-to-end FCoE networks that utilize more than a single switch (that is, multiple FCoE/DCB switch hops from servers to storage). The vendor-centric path Cisco has chosen sacrifices Layer 2 interoperability and possible cost reduction for users in a heterogeneous hardware environment. Instead, Cisco chooses to extend its Fibre Channel switch functionality and protocols from its native Fibre Channel switch products to burden its Ethernet products with the same limitations and lack of interoperability. Fibre Channel networks are inherently single vendor. Therefore, prolonging interoperability with existing Cisco Fibre Channel hardware results in vendor lock-in for every switch in the data center. This is not consistent with customer expectations described in the beginning of this section.

Servers

Enterprise data centers typically pool (share) storage for access over Fibre Channel. For high availability, a server typically has two or more LAN connections and two Fibre Channel connections to provide redundancy. It may have even

more connections for increased performance. Smaller servers, especially blade servers, have few option slots, and the Fibre Channel host bus adapters (HBAs) add noticeably to server cost. Therefore FCoE is a compelling choice for the first network hop in such servers.

HP ProLiant BL G7 and Gen8 blade servers have embedded CNAs, called FlexFabric adapters, which eliminate some hardware costs associated with separate network fabrics. Dual-port, multifunction FlexFabric adapters are both Ethernet and host bus adapters, depending on user configuration. These embedded CNA adapters leave the BladeSystem PCIe mezzanine card slots available for additional storage or data connections.

Direct-attach storage

While Direct Attached Storage (DAS) does not offer an alternative to converged networking, it's important to acknowledge that a great deal of the storage sold today is DAS and that will continue to be the case. HP Virtual Connect direct-attach Fibre Channel for 3PAR Storage Solutions with Flat SAN technology can be a significant way to more effectively use the single hop FCoE already deployed within existing network infrastructures. See the "Using HP Flat SAN technology to enhance single hop FCoE" section later in this paper to learn more.

Storage Area Network

Fibre channel is the storage fabric of choice for most enterprise IT infrastructures. Until now, Fibre channel required an intermediate SAN fabric to create your storage solution. However, this fabric can be expensive and complex. Converged networks using FCoE have the potential to change these requirements in a way that lowers costs and complexity in the IT infrastructure.

iSCSI: SANs using TCP/IP networks

iSCSI infrastructures further demonstrates the desirability of network convergence with respect to cost savings. Two common deployments are the end-to-end iSCSI solution, typically found in entry to midrange network environments, and iSCSI frontends to Fibre Channel storage in large enterprise environments. In the latter scenario, the FCoE benefit is that, from server and SAN perspective, it's all Fibre Channel traffic and doesn't go through translation. The FCoE solution means less overhead and easier debugging, resulting in significant cost savings.

Cluster interconnects

Cluster interconnects are special networks designed to provide low latency (and sometimes high bandwidth) for communication between pieces of an application running on two or more servers. The way application software sends and receives messages typically contributes more to the latency than a small network fabric. As a result, engineering of the end-to-end path is more important than the design detail of a particular switch ASIC. Cluster interconnects are used in supercomputers and historically have been important to the performance of parallel databases. More recently, trading applications in the financial services industry and other high-performance computing (HPC) applications have made extensive use of InfiniBand. Over the next few years, 10 and 40 Gigabit DCB Ethernet will become more common and less expensive. RoCEE NICs (NICs supporting InfiniBand-like interfaces) also will become available. As a result, half of the current cluster interconnect applications may revert to Ethernet, leaving only the most performance sensitive supercomputers and applications to continue running separate cluster networks.

Approaches to converged data center networks

As the industry continues its efforts to converge data center networks, various engineering tradeoffs become necessary. Different approaches may in fact be better for different environments, but in the end the industry needs to focus on a single approach to best serve its customers. Table 1 examines how current protocol candidates compare for converging data center networks in different network environments.

Table 1. Converged Network Candidates

Network environments	FCoE on DCB	iSCSI	InfiniBand	iSCSI on DCB
Directory, security, other, SAN services	Fibre Channel switch software	iSNS (open source)	(minimal)	iSNS (open source)
Flow control (one hop)	Per priority pause (PFC) part of DCB	TCP	per priority buffer credit	PFC
Flow control (end-to-end)	QCN (part of DCB)	TCP	Manual tuning of workload	iSER or QCN
How would a cluster connection be added here?	RoCEE	iWARP (RDMA over TCP)	InfiniBand is the lowest latency, fastest cluster network today	RoCEE
Comments	Best coexistence with and transition plan for Fibre Channel	Success in smaller environments and not enterprise	Success only in low latency and supercomputer environments including storage system internals	Emerging limited vendor support

Convergence strategies

As two of the largest data center infrastructure providers, HP and Cisco have significant impact on the direction of current and future network convergence efforts. Given the nature of its core business, it's understandable that Cisco has a network centric approach. HP is more aligned with the overall business solution in which the software and applications are a central focus. The applications are the tenants of the network infrastructure and HP supports this viewpoint with a broad portfolio of business optimization software. Table 2 compares the contrasting HP and Cisco strategies driven by these different approaches to the data center.

Table 2. This table shows the contrasting data center networking strategies between HP and Cisco.

Category	Cisco	HP
Fundamental strategy for data center infrastructure	Control all data center communications including storage protocols	Develop and provide innovative solutions using industry standard protocols and mechanisms
L2 network design strategy	Hierarchical model with fabric extensions / centralize control	Flatter, less complicated L2 networks / removing the hierarchy where possible
		Leverage intelligent resilient framework (IRF) to facilitate flatter L2 networks
Compute strategy	Control all network end points, including server nodes, from upstream hierarchy similar to traditional network designs.	Increase network end point capability allowing management to occur from multiple levels including software.
		Provide and support compute node connectivity of any type.
Intelligence & management	Push all management to switch devices up the hierarchy	Allow management to occur at node level and promote management distribution
	Remove intelligence from network end points	Develop intelligence at all levels – server, chassis, cabinet, network and storage
Fabric connectivity	Cisco provides solutions in the multi-layer SAN switch segment and strives to provide unique and sometimes proprietary options for an all Cisco infrastructure	remain open and flexible wherever possible to maintain compatibility with existing dc architectures
		support any standards-based upstream network for Ethernet, Fibre Channel and InfiniBand

Category	Cisco	HP
FCoE	Consolidate Fibre Channel traffic with Ethernet / push FCoE for data center	Converge Fibre Channel at c7000 chassis level and in Networking devices where costs are justified and easily defended
	Lead the market in data center (multi-hop) FCoE via proprietary methods	In advance of FCoE standards becoming economically viable and available in end-to-end network environments, HP recommends 3PAR Storage Solutions with Flat SAN technology as the preferred mechanism to provide end to end scalable storage without the complexities of multiple FC/FCoE switch hops.

Cost structure

The real cost structure driver is scale: tens of millions of devices are sold in data centers around the world to achieve economies of scale and to attract competition, resulting in the best pricing. Today, a CNA costs slightly more than a basic NIC; a switch with an FCF costs more than a basic Ethernet switch. In other words, an end-to-end FCoE infrastructure has not yet achieved economies of scale for it to be competitive.

FCoE requires a special network interface card (NIC) that includes most functions of a Fibre Channel interface card commonly known as HBA. This special interface is called CNA. Today most CNAs that exist as PCIe cards, cost as much or more than their respective FC HBAs. Vendors can reduce the cost of adopting FCoE by embedding the CNA on the server motherboard, When the CNA is a single device embedded on the motherboard, it is called a CNA LAN on Motherboard (LOM). Including a CNA LOM on a blade server, as HP does on its latest models, enables Fibre Channel connectivity through FCoE at only slightly higher server cost than running iSCSI over a conventional LOM. This is a compelling cost breakthrough for FCoE/Fibre Channel connectivity.

Congestion and flow control

Data size is an important distinction between storage traffic and other network traffic in the data center. Multiple servers simultaneously accessing large amounts of data on a single storage device can lead to storage network congestion. Managing events when congestion occurs is a basic part of networking and is particularly important for storage networking.

To better understand congestion and flow control, we need to go beyond viewing the network as “speeds and feeds” and examine the pool of buffers. From this perspective, we visualize network congestion not as traffic movement over wires but rather as traffic occupying buffers in a series of devices between source and destination until those buffers run out of space. A single runaway process in a TCP environment overflows buffers, discards packets, and slows down traffic.

In contrast, a single runaway process in a PFC environment pushes packets until a congestion point is found and its buffers are filled. Then the buffers of every device between the sender and the bottleneck fill up, effectively stopping traffic at all those devices. The PFC environment needs a TCP-like mechanism to stop the sender before it floods the buffer pools in so many places. In the current DCB standards, QCN is intended to be that mechanism. As soon as these buffers are filled, other traffic passing through the switches from unrelated sources to unrelated destinations will also stop, because there is no free buffer space for them to occupy. Cisco’s version of PFC uses FCFs in each switch hop instead of QCN. Table 3 compares congestion control available with FCoE, iSCSI, and InfiniBand.

Table 3. This table is a comparison of congestion control technologies

Network type	Single hop flow control	End to end flow control
Ethernet	Drop packets	TCP
InfiniBand	Per priority buffer credit	Tune application (also offers the QCN-like FECN/ BECN mechanism)
Fibre Channel	Credit (BB Credit)	Severely limit number of hops, or tune configuration
DCB	Per priority pause flow control (PFC)	QCN
Cisco’s latest direction for FCoE over CEE	Per priority pause flow control	FCF in each switch hop Same multi-hop limitations as FC SANs

Networks become congested like traffic after an accident. When an auto accident occurs at an intersection in the center of a big city, traffic backs up on the intersection, which in turn backs up traffic in adjacent streets until nothing can move within a full square mile. This is called “congestion spreading.”

The methods for congestion management change with the standards employed. Table 4 shows those changes in different network enviornments and employing different congestion management technologies.

Table 4. This table shows the results of different congestion management methods when end-to-end load overwhelms network capacity

Congestion management technology	Behavior	Technology example showing that behavior can be successfully managed
Drop	Drop packets	TCP
Credit	congestion spreading due to buffers not being granted on links	Severely limit number of hops, or tune configuration
DCB: pause at some priorities, drop at others	congestion spreading on congested priorities due to per priority pause	Tune application or use the QCN mechanism

FCoE and DCB progress and challenges ahead

FCoE at the server edge, from the CNA to the first hop switch with separate Ethernet and Fibre Channel uplinks is widely available in the industry from multiple vendors and should be considered proven. The way bandwidth is allocated between TCP traffic and FCoE traffic varies by vendor.

FCoE from a CNA, through an intermediate Ethernet switch, then to a second hop switch which has separate Ethernet and Fibre Channel uplinks, is slowly becoming available. There are two key technology areas in which industry network vendors promote competing technologies for two-hop FCoE: congestion management, and FCoE Initialization Protocol (FIP) snooping.

In congestion management in two hop FcoE, the way the input of the second switch manages its buffer space determines whether the switch drops TCP packets due to congestion or issues a PFC pause (part of DCB) to lossless traffic like FCoE. While it is always possible to deploy very expensive switches with very deep buffers, most organizations are likely to choose today’s cost effective access layer switches with much smaller buffers inside the switch chip(s). In such designs, much attention goes into optimizing the use of limited buffer space, whether that is fixed assignment of space to ports and classes of service, or quotas for how much each of those can take from a shared pool.

FIP snooping is a feature of the intermediate Ethernet switch in a two hop configuration. In simple terms, FIP snooping is part of an overall security design for first generation FcoE. It is intended to prevent a compromised server on the network from forging FCoE packets. The switch watches FIP packets and only allows FCoE traffic between endpoints that have successfully opened an FCoE connection. FIP snooping is defined in an appendix to the FC-BB-5 standard. Unfortunately, FIP snooping implementations vary so much that almost no two implementations behave identically.

Next generation FCoE is in development in the FC-BB-6 committee at the time of publishing this paper. One of the concepts discussed publicly in FC-BB-6 has been a device called an FDF, which acts as a port expander to a Fibre Channel switch (or Ethernet switch running the Fibre Channel switch software). The FDF will have a much clearer set of responsibilities in the FCoE world than the Ethernet switch with FIP snooping that it will effectively replace.

Native FCoE storage devices are available from some vendors, but not yet deployed widely. We expect that, just as with Fibre Channel devices today, extensive interoperability testing will be needed. Once that testing is complete one expects native FCoE storage devices will work in the tested configurations keeping in line with the FC model of interoperability.

Successful implementation of FCoE through multiple hops of a data center network, or through an arbitrary network topology, requires technology that is still evolving. Interoperability will also be a challenge for at least the next several years. for the next few years, organizations should expect to spend technical resources on configuring and tuning any multi-hop network that carries FCoE traffic.

Cisco has proposed a different approach to FCoE. In the Cisco model Fibre Channel traffic is passed over Ethernet, not as FCoE packets being forwarded at Layer 2 through switches, but instead passing through a FCF at Layer 3 using Fabric Shortest Path First (FSPF) at each hop. This approach calls attention to the following observations:

- Since FCoE was first announced, the Ethernet community has assumed that FCoE packets could move across any Ethernet switch that implemented the DCB features. Cisco's direction represents a fundamental shift away from the Ethernet model and toward the Fibre Channel model in which storage networks always use switches from a single vendor.
- Every switch in the data center requires Fibre Channel switch firmware. The connections between those switches are Virtual Ethernet (VE) Ports. In this Cisco approach, VE Ports emulate an Ethernet port over a non-Fibre Channel link because traffic is running over Ethernet. It is well known that connecting VE ports from two different vendors together forces a Fibre Channel network into a least common denominator mode, which is not useful. FSPF runs as a closed protocol between those switches to build the forwarding tables for Fibre Channel and FCoE traffic.
- Nearly every enterprise customer in the world already has Fibre Channel installed, and thus has chosen between the Cisco and Brocade switch families. Proprietary features of Cisco and Brocade switches prevent them from interoperating well over an Ethernet port. For that reason, most enterprises will run their data center backbone switches on one or the other Fibre Channel code base, and thefore must buy these switches only from the Fibre Channel switch vendor using that code base. The absense of a single Fibre Channel code base makes the approach proprietary and undesirable, limiting customer choice, and inhibiting development of a robust ecosystem. Table 5 expands on the state of convergence standards in the industry.

Table 5. The state of convergence standards

Dependency	Done on paper	Proven in real use	Ready for
Basic FCoE protocol, CNA, FCF, 1-hop connection from CNA to FCF *	Yes	Yes	Mainstream customer use
Flow control (PFC, 802.1Qbb)	Yes	Yes, at small scale	Early adopters
Congestion Management (QCN, 802.1Qau)	Yes	No	Early adopters
FCoE FIP initialization, FIP Snooping, FCoE network security model	Yes. FIP snooping and network security model in an appendix	Multi-vendor interoperability unlikely in this generation	Early adopters
Large data center networks that will accommodate arbitrary FCoE traffic	Still in technical debate	No	Technology experiments
Cisco approach: all switches forward FCoE using FSPF (at Layer 3 as if a Fibre Channel switch)	Yes. inherently single vendor, no congestion management	Yes, at small scale	Deployments no more complicated than existing FC SANs

* Limited support for FCF-per-hop makes it a proprietary, single-vendor approach for the immediate future.

Using HP Flat SAN technology to enhance single hop FCoE

Fibre channel is the storage fabric of choice for most enterprise IT infrastructures. Until now, Fibre channel required an intermediate SAN fabric to create your storage solution. However, this fabric can be expensive, and can result in increased complexity and IT infrastructure costs.

We have improved efficiency of server and storage connectivity with HP Virtual Connect direct-attach Fibre Channel for 3PAR Storage Solutions with Flat SAN technology. You can now connect HP 3PAR Storage Systems directly to the HP Virtual Connect FlexFabric Modules (Figure 3). That eliminates the need for an intermediate SAN switch complex, multi-tier SANs, and excess networking equipment. This innovative solution requires no SAN fabric licenses. In an existing fabric-attach environment, you can use the 3PAR Storage Solutions with Flat SAN technology to direct-attach and fabric-attach storage simultaneously. More information is available in "HP Virtual Connect direct-attach Fibre Channel for HP 3PAR Storage Systems solutions brief" at http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA4-1557ENW.pdf

Figure 3: HP Virtual Connect FlexFabric direct-attach with FC based HP 3PAR Flat SAN technology Systems

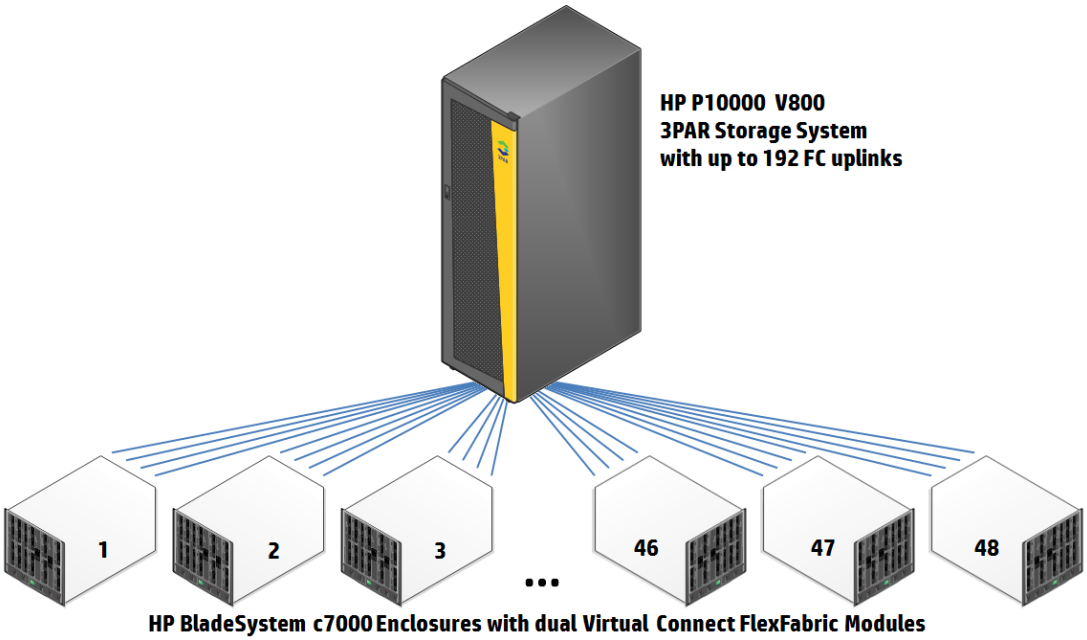

Evolution of storage

It seems inevitable that the storage networking now accomplished with Fibre Channel will transition onto an Ethernet-based converged standard such as FCoE throughout most of the data center. In the past, several attempts at accelerating this transition assumed that a single and fairly rapid technology replacement would occur. But history proves that pervasive changes occur rarely and slowly. Perhaps it is useful to consider a spectrum of Ethernet connected storage alternatives.

Ethernet connected NAS appliances serving files have been emerging for decades and have a respectable share of the storage market; likewise in very high-end supercomputers a large parallel file system (such as HP IBRIX) runs on some of the nodes, with applications accessing data as files over the network.

To reduce cost and make latency/throughput predictable, applications such as Microsoft Exchange are evolving away from SAN storage and back to DAS. Perhaps you can think of these cases as an application-level protocol on the network, not as a block storage protocol.

iSCSI has been very successful at providing block storage in smaller environments. The HP Left Hand Networks clustered iSCSI products are an example.

FCoE is a very compelling edge connection for Fibre Channel storage networks, especially in blade servers with CNA LOMs. For at least a decade to come, customers who already have Fibre Channel can continue using the hybrid FCoE edge, Fibre Channel core design.

iSCSI over DCB offers many of the advantages of the FCoE/Fibre Channel combination. Using open source Internet iStorage Name Service Server (iSNS) instead of the higher priced Fibre Channel switch software from the major vendors will potentially keep costs much lower and eliminate the need for translation in Fibre Channel Forwarders. The industry may find a way to build an iSCSI/FCoE hybrid with common services across the two, but it is too soon to know whether such a product will emerge and if so, whether customers will embrace it.

Some organizations are late adopters and will simply stay with Fibre Channel for a long time. There is a lot to be said for not spending resources on a big change until it is truly needed. The IT industry and customer base can only support one technology as the primary storage connection in the data center. The availability of a wide range of products, well developed expertise within the given technologies, and the maturity of the technology all follow the volume generated in the marketplace. We can't be certain how the available choices will fare. In the meantime it is important not to push one-size-fits-all, but rather to find the best storage connectivity for each organization.

Conclusion

Server virtualization is creating significant change, driving data centers to pool larger numbers of servers. There is a need to replace the rigid compartmentalized networks with a more flexible model which still contains the impact of a single runaway process or switch can have on the network. The winning approaches have not yet emerged, or more precisely, every vendor is selling you a different “winning approach” that addresses some but not all of these issues.

As a major vendor behind FCoE, Cisco is driving efforts for a data center switch refresh. The aim is to replace aging technology with their new generation switches, using FCoE requirements as the reason to act. A year ago, Cisco changed plans from traffic as FCoE over a large Layer 2 DCB fabric with congestion management (using Congestion Notification to traffic from FCF to FCF as switch hops are traversed). At best, this makes the industry direction for multi-hop FCoE unclear, and at worst indicates that end-to-end FCoE will only work on single vendor networks. Such uncertainly discourages long term investments based on assumptions about FCoE’s requirements.

Our advice on native, multi-hop FCoE across the data center includes FCoE as a strong server edge technology for native Fibre Channel networks, with 1-hop FCoE proven and 2-hop FCoE viable in selected cases. End-to-end FCoE is viable only in single vendor proprietary networks for the foreseeable future. Any deployment of end-to-end FCoE should be carefully scrutinized for cost, benefit, risk, and ROI. Consider adopting technologies like the FC-based HP 3PAR Flat SAN technology Systems that operate within the proven single hop environment.

This is a good time to explore storage connection alternatives and begin to estimate what percent of your storage needs is best met by block storage on a SAN (Fibre Channel, FCoE, iSCSI), what percent is best met by file storage (NAS, and at what scale), and what percent is best met by DAS. These are some of the other planning considerations:

- On what scale do you implement migration strategies (like vMotion or others) requiring that MAC and IP addresses stay intact?
- Do you intend to move from the classic oversubscribed hierarchical network to a flatter network topology. Hierarchical networks can cause unacceptable oversubscription on both data and storage networks, imposing more upward hops on both kinds of traffic.?
- What choices should you make between the various hypervisor and network vendors as next generation data center networks emerge?

You can’t ignore escalating business requirements, but you can balance those drivers with a measured approach to the adoption of new network technologies. Evaluating new technologies based on how they preserve the value and remaining lifecycle of existing network infrastructures should be a consideration. It may also be worth waiting to embrace a new end-to-end network strategy until there is some consensus about the DCB standards among network hardware vendors.

For more information

Visit the URLs listed below if you need additional information.

Resource description	Web address
Comparison of HP BladeSystem servers with Virtual Connect to Cisco UCS technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c03094466/c03094466.pdf
Converged networks with Fibre Channel over Ethernet and Data Center Bridging technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01681871/c01681871.pdf
HP Virtual Connect Direct-Attach Fibre Channel for HP 3PAR Storage Systems solutions brief	http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA4-1557ENW.pdf
HP FlexFabric Networks web portal	http://h17007.www1.hp.com/us/en/solutions/datacenter/index.aspx
OpenFlow - Software-Defined Network	http://h17007.www1.hp.com/us/en/solutions/technology/openflow/index.aspx?jumpid=reg_r1002_usen

Get connected

hp.com/go/getconnected

Current HP driver, support, and security alerts delivered directly to your desktop

TC1207917, Created July 2012